

MANET MAIL ABOUT EVALUATION

Temporally-Ordered Routing Algorithm (TORA) - Vincent Park

See <http://tonnant.itd.nrl.navy.mil/tora/tora.html> for more information on TORA.

TORA uses a metric referred to as the "height" of the node to assign a direction to links for forwarding packets to a given destination. The node heights can be totally ordered lexicographically, and thus define a directed acyclic graph rooted at the destination.

There are three functions: creating routes, maintaining routes and erasing routes.

Creating routes: Creating routes is performed on demand using a query/replay process.

Maintaining routes: When a node loses its last downstream link the algorithm reorients the directed acyclic graph such that all downstream paths lead to the destination.

Reoptimization of routes: TORA does not compute the shortest path: paths may be suboptimal. It starts close to optimal and tends to "loosen", as it reacts to topological changes. A secondary mechanism, not tied to the rate of topological change, is used to reoptimize routes.

Partition detection and erasing routes: Partitions are detected when a 'reversal' reaches a node with no downstream links and all of its neighbors have the same 'reflected reference level,' which it previously defined. A node that detects a partition initiates the process of erasing the invalid routes.

Simulation:

Protocol comparison: Performance of TORA was compared to Ideal Link-State (ILS) and pure flooding. Since TORA often provides multiple downstream routes, a next-hop forwarding decision is required. Two different forwarding policies were evaluated: TORA - for each packet randomly(based on uniform distribution) select one of the downstream neighbors to forward the packet to, and TORA LN - forward all packets to the "lowest" downstream neighbor.

Simulation description: Simulations used fixed network topologies with the ability to fail and recover links based on an exponentially distributed time intervals. This was an adjustable parameter used to vary the rate of topological change. Other parameters used to vary average network connectivity and message traffic load. Multiple baseline topologies were used to evaluate the effect of network size on routing performance. Performance comparison was based on measure of control and data traffic, end-to-end message packet delay and message packet throughput. As rate of topological change was increased, the control overhead for ILS increased significantly faster than for TORA. Excessive ILS overhead caused network congestion, resulting in longer end-to-end message delay. TORA outperformed ILS (in terms of bandwidth utilization and end-to-end delay) under conditions of high rates of topological change. As network size was increased, TORA outperformed ILS at lower rates of change. The throughput

statistics provided little insight into the difference between TORA and ILS, and thus were not presented.

Traffic distribution was chosen to be the WORST case for TORA. Every node generated traffic (at exponentially distributed inter-arrival times) destined for every other node in the network.

Summary of results: TORA performs better (than ILS) as rate of topological change is increased and as network size is increased. Network connectivity did not significantly affect relative performance of protocols.

In cases where the network is very static, it is better to use ILS. Otherwise, use TORA.

Question: Does TORA always converge? Answer: TORA converges probabilistically with time. However, an example has been constructed, which shows that under certain conditions TORA can exhibit oscillatory behavior and need not converge within a finite time. The example is dependent on a specific topology and specific timing of events (packet transmissions), which makes it highly unlikely for the behavior to continue for multiple cycles. Vincent and Scott stated that there is a solution which guarantees convergence. An unscalable solution would be to only build routes from the destinations.

Question: Link up/down simulation was used to model motion. Isn't this a problem?

Answer: No - This is an acceptable model for simulation of protocols that do not benefit from spatial/time correlation like TORA.

Simulating Wireless Nets

Chip Elliott (celliott@bbn.com)
Fri, 25 Jul 1997 08:20:19 -0400
To: manet@itd.nrl.navy.mil
From: Chip Elliott <celliott@bbn.com>
Subject: Simulating Wireless Nets

Howdy,

Bevan, thanks very much for that bibliography. You guys seem to be doing good work with your ad hoc routing.

We at BBN are doing a fair amount of simulation and modeling for the NTDR ad hoc network. Here are the things that we have found useful thus far:

- a) Pathloss model for a given terrain (eg rolling hills with foliage, desert mountains, etc; we try to use real places rather than dreamed up models)
- b) Initial laydown of nodes on terrain
- c) Mobility scenario, which includes:
 - 1. trajectories of nodes
 - 2. nodes powering on/off over time
 - 3. various types of jamming
- d) Traffic scenario

e) Exact modeling of the protocols and algorithms (in our case, same code in OPNET model, simulator, and radios)

In my opinion, models must have reasonable fidelity in the following areas or they will give very misleading results:

- Pathloss (free space gives VERY unrealistic results)
- Traffic (most traffic is highly correlated, rather than uniform distribution, and this affects interference)
- Trajectories (the usual 'brownian' gives very optimistic results, as no one really goes anywhere)

Naturally the desire for realism conflicts horribly with the desire to run simulations of large (400 node) nets.

We have a reasonably elaborate system of "scenarios" which feed into both our OPNET model and our 50-node testbed, so we can run experiments repeatedly and measure the results as we tune things. A scenario consists of a selection of values for items a-d above.

Another interesting set of questions is... what do you want to MEASURE...?

Simulation parameters for comparative results?

Bevan N. Das (das@grinch.csl.uiuc.edu)
Thu, 24 Jul 1997 10:35:19 -0500

I've had one nagging thought about the simulation results in the different papers on routing in ad hoc networks (including our own). As Park and Corson noted in their TORA paper in Infocom '97: "what is important is the protocol's average performance which is only obtainable from simulation." However, there is no common ground to compare the simulation results from different authors; it's like comparing apples and oranges.

Has there been discussion or conclusions regarding parameters for simulation scenarios? Some possibilities I can think of:

- network size
- network topology, such as average degree/connectivity of nodes, also any specific topologies
- physical considerations, such as area covered, transmission range, received power, bandwidth constraints.
- mobility model, such as the rate of change, the type of changes modelled--i.e., independent link changes vs. movement of nodes (so that link changes are NOT independent), the choice of static topology with random failures vs. mobile topology with random movement

In summary, what would make a "good" simulation scenario that would be a basis for comparison? (in contrast to "something that we can program").

Re: Simulation parameters for comparative results?

Noel Chiappa (jnc@ginger.lcs.mit.edu)
Fri, 25 Jul 97 10:22:40 -0400

From: "Bevan N. Das" <das@grinch.csl.uiuc.edu>

"what is important is the protocol's average performance which is only obtainable from simulation." However, there is no common ground to compare the simulation results from different authors

Indeed! It can be hard to compare the work done on different algorithms. (I prefer to think, and talk, about 'algorithms', rather than 'protocols', for a number of reasons, such as encouraging people to draw lessons that are more broadly applicable than a specific protocol.) I'll comment on your specific suggestion below, but before I do that, here's one thing to keep in mind - and I apologize if you're already well aware of this.

To me, one of the most important results to convey is something about the *shape* of the stabilization time (i.e. the time it takes to produce stable, fully adapted routes after a topology change) curve, not just a single average. For example, a routing algorithm that stabilizes in time T 98% of the time, and $100T$ the other 2%, is going to show a nice "average" of about $3T$ - but would not really be acceptable in service.

I don't know whether simply giving a standard deviation would be good enough, or whether you'd need more of a graphical result - but a single number is *definitely* not good enough.

The way I would suggest measuring it is to take a given 'sample' topology, and measure the stabilization times for each potential topology change. (It might be interesting to run the simulation a number of times for the *same* topology change, and see how much the stabilization time varies. If you get a lot of divergence there, that would also be interesting to know. This assumes that you have some randomization in your model, to account for real world things like the effects of competition from user data for access to transmission facilities, etc.) After producing a curve, try it for a different topology, and see if you get a very different curve. If you do, that's also an interesting thing to know.

parameters for simulation scenarios ...
-network size

A good one. What is especially interesting to know is if the *shape* of the curve stays the same, offset by some factor (perhaps a polynomial) of the network size, or if increasing size causes negative changes in the *shape* of the curve, e.g. increasing the length of the tail.

-network topology, such as average degree/connectivity of nodes

Graph theory says that as you increase the degree, average path length in a network should decline. This is more likely to have radical effects in routing algorithms in which the path selection part of the algorithm is distributed (e.g. distance vector, path vector), as opposed to algorithms in which path selection is done in parallel, locally (e.g. link state). Probably less interesting than the first, though - but the same comments about seeing if the shape changes apply.

-physical considerations, such as area covered, transmission range, received power, bandwidth constraints.

I would think some of these are less likely to affect the results than others, but it would be interesting to try a few to see. Are you considering them only from the point of local, mobile, networks, or are you thinking of results that have broader applicability? E.g. bandwidth issues (e.g. from interference with user data traffic, as well as actual transmission delays for routing data, for those classes of algorithms that have to ship lots of data around when a change with significant impact happens, such as most destination vector algorithms) are probably less significant in the network as a whole than in over-the-air networks of limited bandwidth. Conversely, physical end-end speed-of-light transmission delays are more important in continental and global networks than they would be in local networks.

(I am particularly interested in SOL issues as they are **increasingly** important as other speeds (processor, memory and link) decline, while SOL delays are constant, and will, I believe, come to be an important factor in the "equation" of stabilization time for the network as a whole.)

- mobility model, such as the rate of change

This is less interesting to model, as if you know the shape of the stabilization curve, you can more or less predict where you are going to get into trouble - when the average stabilization time is more than the mean time between topology changes - at which point the routing starts to chase its own tail.

the choice of static topology with random failures vs. mobile topology with random movement

I'm not sure this is really a difference, at least for anything you're likely to program. The latter just looks like links that have been down "forever" coming up - and as far as most path selection algorithms are concerned, a topology change is a TC is a TC, as far as what the algorithm has to do to respond.

If you had a large enough network that you needed to use aggregation to limit the routing overhead, you have an interesting thing where after enough topology changes, you might have to re-do abstraction boundaries on the fly, and rename nodes, but I don't think you're likely to simulate this.

In summary, what would make a "good" simulation scenario that would be a basis for comparison?

I think you need to be able to either i) feed in an actual topology, or ii) have it simulate one. (In fact, maybe the topology generator could be a separate program, so the actual simulator doesn't have to know what the source is.) I have my suspicions that "randomly generated" topologies have different characteristics than real networks. For one, I think there is quite likely a different "spacing" (this is a graph theory term I invented to describe the **kind** of connectivity a node has, local or distant, as opposed to just the **degree** - but it turns out to be much the same as "diameter", for a given graph, when you think about it) in real networks and in purely random graphs. In fact, there's an interesting piece of research right there: try some real networks, and random ones, and see if they **do** differ in their actual results, so in future we'll know whether we have to use real connectivity data, or can generate random ones.

I also think it's less critical that you generate numerical data that can be compared with other models (e.g. "algorithm X with Y nodes stabilizes in Z seconds"), than you do internal comparisons, with curve outputs which can be interpreted directly (e.g. "algorithm X produces this kind of curve, and it stays the same shape as the network gets larger, and algorithm Y produces this shape curve").

This is all kind of a gloss on a very complex topic, actually, but hopefully these quick jottings will be of some use.

common performance measures

Kimberly King (kimberly@albers.tieo.saic.com)

Tue, 29 Sep 1998 09:34:15 -0400

I agree we need to find a common basis for measuring protocol performance.

However, if one represents this common basis in terms of physical movement of nodes, then the actual network context is highly dependent upon physical and link-layer assumptions. By physical movement of nodes I mean specifying node A is at (x,y) at time t0. By network context I mean

- # of nodes
- average number of neighbors
- average link lifetime
- # of times the network becomes partitioned
- average length of network partitions.

If comparisons are based on abstract network characteristics, then one avoids the issues of dependencies upon the physical and link-layers. As a mathematician, I believe better representations for mobility exist than using simulated physical movement on an (x,y) plane.

It is also important to determine the stability of performance results. If only a few fixed models using the physical movement notion is adopted then it is difficult to assess how that particular topology influences performance. Whatever representation is adopted, it is helpful to be able to see if the same results are obtained given the same network context.

In addition to mobility representation, I think we need to separate representation of data traffic. Clearly, the same network context with different traffic characteristics affects routing performance.

Scenarios: mobility

Miguel Sanchez (misan@acm.org)

Tue, 29 Sep 1998 09:56:40 -0000

As most of you, I've been reading with interest the recent article from CMU researchers (it is shorter than the five author's names :-). I really enjoyed it. I only want to point out a fact that probably most of you have think about. Mobility models are important to be agreed on if we want to be able to compare simulation results (not only routing protocols).

In fact, physical and data-link layer offer a variety enough to change results greatly. Transport protocol, of course, also imposes some restrictions on wireless links. BTW, I recommend you an article regarding TCP enhances to be used over wireless links:

* H. Balakrishnan, V.N. Padmanabhan, S. Seshan, R. H. Katz "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links" IEEE/ACM Transactions on networking Vol.5.NO.6 December 1997 Pags 756-769

Regarding to mobility, it is clear for me that the more randomness of movement is, probably the worst case of motion to the routing protocols. Therefore, in order to test a routing protocol behavior I would recommend a random (may I say "pattern"?). The motion model of the CMU article is similar to the movement of several shopping baskets inside a supermarket (that's appreciation is mine, the authors may not agree). I think it is OK due to its high randomness.

However, real world motion scenarios would not present, likely, an excessive randomness but, more probably, a less chaotic motion. This includes motion restrictions due to the environment, power conservation issues (machine motion may be expensive in battery consumption), and, of course the application level (the ultimate purpose of this device). On the other hand, Autonomous systems (robots) would not exhibit the same kind of motion than human held devices.

Putting all this short: A common motion model is needed. (I think a highly random one for protocol stress testing and a second one, less random, for performance measurements).

Re: Scenarios, Assumptions, and Evaluations

Joe Macker (macker@itd.nrl.navy.mil)
Mon, 28 Sep 1998 14:08:35 -0400

At 03:41 PM 9/25/98 -0400, Strater, Jay wrote:

>

>Given the interest in defining net scenarios and traffic assumptions and in
>determining the conditions/assumptions under which different protocols
>perform best, we'd be happy to share our conditions/assumptions and
>performance results in the next MANET meeting in December. Currently, we
>are in the middle of our simulations so we don't have adequate information
>to share with the group. By the next meeting, however, we expect to have a
>fairly large set of results.

Yes, we do want input on this and recommendations on scenario sets, etc. Input is, of course, welcome in December.

While next meeting is o.k., input you can discuss now on the mailing list would be appreciated. Even without results; your group size, mobility models, assumptions, traffic scenarios under consideration would be of interest. Any early strong opinions you have formed on effectively covering the broad range of performance considerations would be welcomed.

Scenarios, Assumptions, and Evaluations

Strater, Jay (jstrater@mitre.org)
Fri, 25 Sep 1998 15:41:08 -0400

I've been late in reading my MANET mail but now that I've finally made it through I realize that I should share with you the state of MITRE evaluation activities.

As you and a few others may know, several colleagues and I at MITRE have been involved in conducting evaluations (via OPNET simulation) of some advanced routing protocols (MANET and DARPA) for tactical military ad hoc nets. The purpose for the evaluations is to determine the appropriateness of the protocols for different tactical conditions. In the evaluations we are simulating protocol performances for a variety of tactical laydown/movement/terrain scenarios, traffic services, net sizes, and loading conditions. In particular, we are simulating "external" performance metrics (such as delay and throughput) and "internal" performance metrics (such as information and control efficiency and route lengths). We've chosen 3 scenarios that have varying levels of connectivity and link fluctuation characteristics. We've chosen a mix of traffic types that range from reliable to time sensitive traffic types. And we're evaluating net sizes of 10, 20, and 40 node nets. In each case, we've varied loading to the point where end-to-end delay becomes excessive (to determine the limit of the protocol and net). Although our scenarios and traffic assumptions are military in nature, we've characterized them in a manner that may be comparable to commercial scenarios and traffic assumptions.

Given the interest in defining net scenarios and traffic assumptions and in determining the conditions/assumptions under which different protocols perform best, we'd be happy to share our conditions/assumptions and performance results in the next MANET meeting in December. Currently, we are in the middle of our simulations so we don't have adequate information to share with the group. By the next meeting, however, we expect to have a fairly large set of results.

Re: Simulations & Comparisons

M. Scott Corson (corson@glue.umd.edu)
Mon, 21 Sep 1998 23:01:50 -0400 (EDT)

> Routing Power rates B more highly than A due to B's relatively lesser delay.

On second reading, RP is not all that useful either...sigh... ;-(, as it tries to quantify a trade-off which is inherently application-specific.

Being in a good throughput-delay trade-off cannot necessarily be indicated by their ratio. These two algorithms with T/D metrics

100kbps / 1 sec = 100 (measured in Rout. Power)

10kbps / .1 sec = 100 (measured in Rout. Power)

give the same RP, but either might be better-suited for a given application depending on its throughput and delay requirements.

So...as a group, what do we want to target in terms of initial performance trade-offs? Do we want to primarily maximize network-layer throughput whenever possible (to minimize the need for end-to-end retransmissions), and secondly minimize latency? This would seem to be better-suited for supporting file transfer and the like. Is that a reasonable near-term goal? If not...what? Should such be added to the applicability statement for each protocol indicating perhaps the types of applications that might be favored by a given approach?

There was significant discussion concerning an applicability statement during the meeting, but very little since I posted that initial cut. Is everyone happy with the version I sent a week or so ago?

Simulations & Comparisons

M. Scott Corson (corson@glue.umd.edu)

Mon, 21 Sep 1998 18:08:23 -0400 (EDT)

Thanks for contributing the ns-2 manet simulation models to the group, and for taking a first cut at simulating several of the proposed manet protocols---it took a lot of effort.

Regarding your mobicom paper...

It appears likely that there are problems with the IMEP/TORA simulation models, but that is somewhat understandable as the respective I-D's were and are still incomplete.

In terms of the performance metrics comparing the algorithms, after reading the paper, I find myself wondering about two metrics:

1) Route acquisition latency

2) Data delivery latency

The issue of route acquisition latency has already been raised on the list (and subsequently incorporated into the MANET routing protocol performance issues draft) as a potential negative consequence of on-demand (reactive) routing algorithms. This metric would be very useful in comparing various reactive routing algorithms, as well as comparing reactive and proactive approaches.

Data delivery latency is one of the most important performance metrics as it applies to all routing algorithms, and it is not included in the paper. Along with data throughput (which you have computed in percentage form in the packet delivery ratio), it is fundamental to comparing routing algorithm performance. Without this metric, I do not know how to realistically assess routing algorithm performance.

For the working group effort, I think a protocol performance comparison should include these metrics.

An issue for the group...

Once end-to-end data throughput and delay are known, it is still an interesting problem to assess which protocol is better for some environment. For instance, if one has two protocols with the following performances:

Protocol	Throughput	Avg. Pkt. Delay
A	95%	10 sec.
B	90%	9 sec.

Which is better? What else (other than (1) above) should we consider?

The metric of "routing power" has been devised as a means of quantifying the Throughput-Delay relationship by creating a ratio of the two.

Routing Power = Throughput / Delay

It normalizes throughput with respect to delay. Here,

Protocol	Throughput	Avg. Pkt. Delay	Routing Power
A	95%	10 sec.	9.5%/sec.
B	90%	9 sec.	10.0%/sec.

Routing Power rates B more highly than A due to B's relatively lesser delay.

Routing Power is not the ultimate measure, but it should be considered along with other factors. For example, from the perspective of assessing a routing protocol's effect on transport layer performance, routing power probably underpenalizes the effect of missed network-layer throughput on reliable end-to-end delay (e.g. TCP), as missing packets must be retransmitted end-to-end at the higher layer. But it's probably a reasonable measure for predicting a protocol's effect on best-effort, end-to-end data delivery performance.

Date: Mon, 6 Apr 1998 11:50:36 -0400 (EDT)
 From: "M. Scott Corson" <corson@glue.umd.edu>
 To: MANET <manet@itd.nrl.navy.mil>
 Subject: Issues...

1) Simulation Models/Contexts for Protocol Comparison

It is desirable to understand how the various proposed routing protocols compare with one another in terms of performance, and to know in which networking contexts certain approaches are preferred. To do so, we first need to identify what are the networking contexts of interest to this group. Possible networking contexts are:

- * small-scale, ad hoc meeting
- * medium-sized, conference context
- * large-scale, military context

Can people suggest others?

We also need to choose suitable values/models for each of the following parameters to define a context. Some relevant parameters are:

- * network size
- * number of physical-layer technologies
- * available bandwidth/power for each technology
- * nodal degree per technology
- * fraction of unidirectional links per technology
- * router motion models
- * traffic models/patterns
- * fraction of sleeping routers

Comments? Have any significant characteristics been overlooked?

Ideally, it would be nice if everyone in the working group used the same simulation package. Then individual protocol developers could code their protocols and distribute the models to the rest of the group for mutual comparison. Short of this, significant recoding of protocols will be required for each simulation package. Different participants seem to be using different packages: Opnet, NS2 and Maisie are the three of which I'm aware. Each has strengths and weaknesses. Maisie and Opnet lack detailed upper level protocol support. Opnet also lacks a good wireless mobility/topology model for MANETs, and the simulations which have used it

tend to use a fixed topology. NS2 has great upper level protocol support, and may now have a good wireless mobility/topology model depending on what Dave has developed. This makes usage of NS2 very promising as its missing ingredient--the network layer routing code--is what each protocol proponent can develop and distribute to the others.

Dave: You mentioned in the meeting that this might be available by Summer. Why not now? ;-) If you've done what the group requires, it would be great to get everyone using it ASAP.

Date: Fri, 17 Apr 1998 09:41:17 -0500
 From: Bevan Das <bevandas@lucent.com>
 To: "M. Scott Corson" <corson@glue.umd.edu>
 Subject: Re: Issues...

One more parameter to include: the message size. Both the absolute size (e.g. 512 bytes, 53 bytes) and the size relative to the network size. I.e., are messages only $O(\log n)$ size, and hence can contain information about a few nodes? or are messages large enough to contain entire routes? Also, what measures do we use to compare simulation results? Possibilities:

- * throughput
- * delay
- * overhead
- * length of routes compared to shortest paths
- * the ratio of unsuccessful route queries to overall route queries

These questions came up in our work on spine routing in ad hoc networks,

Subject: Re: Issues...
 Date: Sun, 19 Apr 1998 09:40:37 -0700
 From: Bora Akyol <akyol@bbn.com>

Does your code also include effects such as Rayleigh and Log-normal fading? That would be excellent. Is your model based on the Jakes model for propagation?

To: "M. Scott Corson" <corson@glue.umd.edu>
 From: Dave Johnson <dbj@cs.cmu.edu>
 Subject: Re: Issues...
 Date: Sun, 19 Apr 1998 00:53:46 -0400

Sorry for the delay in replying -- we're planning to release our simulation code for NS2 soon, but right now, we're working on a number of things as fast as we can, one of which is documenting the simulator so others besides us can install and use it, and another is actually using the simulator :-). We'll release the code as soon as we can, which will be sometime in early summer, after a number of other deadlines (and classes for this semester) are over.

For those not at the MANET meeting in Los Angeles, our extensions to NS2 provide a detailed simulation of the physical and link layer behavior of a wireless network, and simulate movement of nodes within the network. At the physical layer, we provide a realistic simulation of factors such as free space and ground reflection propagation, propagation delay, transmission power, antenna gain, capture, and receiver sensitivity. At the link layer, we simulate the

complete IEEE 802.11 wireless LAN DCF MAC protocol. The simulator allows programmable node mobility and communication patterns and operates on a terrain defined by a loadable digital elevation map.

Simulation environments

Christian Joensuu FOA 72 (chj@lin.foa.se)
Wed, 10 Feb 1999 15:38:34 +0100

I'm curious about what simulation environments you use working on issues regarding mobile ad hoc networks, routing protocols, and mobility models?

We've been looking at the freely available PARSEC environment for our project and we're hoping to lift our digital terrain based mobility model into such environments and evaluating both access and routing protocols of our own in such an environment for our work in the MANET area.

Do you have any suggestions based on your own MANET simulations experiences?

Any comments are greatly appreciated.

PS Please direct e-mail to me directly since perhaps not all on the list might be interested in your replies. (If I get any...)

Joe Macker (macker@itd.nrl.navy.mil)
Wed, 10 Feb 1999 13:46:55 -0500

Work has been done in a number of environments (e.g., MAISSIE, OPNET, and ns2).
BTW, mobility model input would be useful.

Of particular note, the group made some initial progress towards some common models in the freely available ns2 environment. This work is based upon extensions done at CMU and also there has been work in ns2 done at Sun.

<http://www.monarch.cs.cmu.edu/cmu-ns.html> (CMU ns extensions)

Not sure of the publicly available status of other work or updated information. Anyone?

Lee Chee-Jwai (lcheejwa@dso.org.sg)
Thu, 11 Feb 1999 09:08:24 +0800

Has any comparison of the various simulation tools (MAISIE, OPNET and ns2) been done? What are their respective strengths and weaknesses? Thanks.

Scenarios for simulations

Thomas Lofgren (lofgren@sics.se)
11 Feb 1999 15:35:56 +0100

I'm currently in the process of setting up some interesting mobility scenarios for testing some of the proposed routing algorithms. I'm right now considering to do simulations on (at least) the following two scenarios:

- (1) a high-density clustered environment, such as people carrying around laptops at a fair, or moving between lecture halls. This would involve having one or more completely fixed nodes.
- (2) a disaster relief operation, where some semi-fixed units can provide routing at strategic places, but where "end nodes" may be much more mobile.

For the second one, I will contact people who work in this field to listen to their input, but I will gladly accept any input on both of these scenarios. Is there a reason not to do them? Is someone working on similar or related scenarios? What to think about when designing the movement patterns, etc. Also, any ideas for other scenarios will be considered as well.

Rupert Goodwins (Rupert_Goodwins@zd.com)
Thu, 11 Feb 1999 17:53:50 +0000

Another scenario I'd like to see modelled is that of a domestic metanetwork. A typical example of this would be a residential street where a selection of houses have (say) Bluetooth or other in-house LANs, and decide to create a radio link between themselves. As this will be inherently unreliable -- stuff will appear and disappear as it's turned on or off -- it resembles a classic ad-hoc mobile network more than it does a traditional wired fixed network.

I can see this sort of thing assuming considerable commercial importance, once the density of users gets high enough...

Sanket S. Nesargi (sanket@utdallas.edu)
Thu, 11 Feb 1999 17:25:27 -0600

I was just wondering, if the Manhattan model used for traditional cellular networks made sense here. The nodes would be modeled as moving along a grid representing streets in a downtown and can turn only at right angles. This could lead to the breakage of links, if the turns were around buildings leading to loss of sight of the antennas and due to absorption/multi-path fading due to buildings. Such a grid based model could also make sense, say for a group of fire-fighters equipped with low-power communication devices (forming an ad-hoc network) when they may be trying to fight a fire in an urban setting: both inside and outside buildings.

Also, Thomas had suggested scenarios for "a disaster relief operation, where some semi-fixed units can provide routing at strategic places, but where "end nodes" may be much more mobile." This appears to be more of an application of virtual cellular networks, where there are mobile base stations, and more mobile end nodes. In such a model we have the added luxury of the mobile base stations having more resources than the mobile terminals. However, issues relevant to mobility and resultant changes in topology still apply in such a model. Also, if connection oriented communication is desired, as in cellular networks, this virtual cellular model poses some interesting challenges. A description of this system model can be found in our Infocom '99 paper "Distributed Wireless Channel Allocation in Networks with Mobile Base Stations" available at <http://www.utdallas.edu/~ravip/paper.list.html>.

C.-K. Toh (cktoh@ee.gatech.edu)
 Thu, 11 Feb 1999 20:38:38 -0500

> How about the Paris/DC Lafayette model where core streets are radial and travelers are forced to occasionally turn down chaotic pseudo-grid side streets that may dead-end at a moments notice... Thus causing them to call for assistance quite frequently over their mobile ad hoc network...;->

>

> -joe

I think using a confined migration model is less likely to be appealing. Ad hoc networks are less rigid than cellular in terms of structure and hence it should not be restricted to specific environments.

Joe Macker (macker@itd.nrl.navy.mil)
 Fri, 12 Feb 1999 10:18:49 -0500

I agree with you, I was trying to spread a little humor based upon the history of poor city planning (how chaotically-structured some city side streets can become). I should have explicitly indicated a sarcasm tag in my posting.

The point on overly confined models is well taken and I wholeheartedly agree. Another additional issue is if one does use a highly detailed confined urban model, reflection/fading/blockage/topography and other considerations also affect wireless topology..its not just a flat grid location/range problem. The detailed 3D environment plays a role.

While these issues are important for real system deployment and planning, I believe more open mobility and traffic models help determine the generic value of a mobile routing algorithm(s) across a range of conditions. onethless, its nice to get some data and discussion on particular scenario conditions that might represent urban usage and also real world movement patterns in other than urban environments.

Ravi Prakash (ravip@utdallas.edu)
 Fri, 12 Feb 1999 10:21:56 -0600 (CST)

I agree with Joe and Sanket. And I did see the sarcasm tag! So, no problem there.

A flat grid may be alright to obtain generic data, and it may satisfy several of us (including me on several occasions) who model mobility as a random movement of nodes in this flat region. However, in reality, mobility patterns may be quite different and we do need to consider reality sooner or later. :-)

Also, I think Sanket was trying to raise issues of multi-path fading, blocking, absorption when he referred to the Manhattan model, which BTW is used in the cellular telephony community to mean urban settings with lots of buildings separated by streets running parallel and perpendicular to each other.

If we consider an ad-hoc network to be composed of devices carried by pedestrians in the downtown area we need to consider blocking/fading/reflection, etc. One may be surprised by the frequency with which links appear and disappear. Signals don't fade gracefully. There are sudden

drops. This has the potential to affect the performance of all the routing protocols being discussed by MANET.

For example, if it turns out that routes between pairs of nodes are short-lived, we may have to take a hard look at the caching policies of routing protocols like AODV and DSR. Even if they are not short-lived, the duration an entry should stay in the cache before being purged should be influenced by the terrain. Any protocol that assumes that paths remain fairly stable may have poor performance in such a situation.

What if one wants to use some resource reservation/admission control based policy for QoS guarantees? If the paths change rapidly due to the peculiarities of the environment, the game may not be worth the candle: the effort of setting up reservations (as in RSVP, for example) may be more than the subsequent performance gains.

Also, if we consider in-building ad-hoc networks, the flat grid model is simply not acceptable. The architects may place "stuff" that is aesthetically pleasing, but plays havoc with the propagation characteristics. Apparently, the nice looking tall glass windows at the Denver airport absorb signals like a sponge. I'm told it was a nightmare for the service providers.

Re: Node Movement Models in Ad hoc

George Aggelou (G.Aggelou@ee.surrey.ac.uk)
Thu, 15 Jul 1999 11:11:29 +0100

Have a look at Miguel's home page, at
<http://www.disca.upv.es/misan/mobmodel.htm>, for a discriptive reference on various mobility models...

On my simulations I am using a moving pattern similar to a "realistic" model where mobiles roam freely throughout the coverage area according to the following moving algorithm:

At any point in time, mobile movement is characterized by two parameters, its velocity (with value v and direction f) and its current position(x, y). In this model, then new moving direction at time $(t + 1)$ is computed based in the previous direction at time t plus a steering angle, which is generated randomly by a normal distribution process with some mean, a , and deviation L degrees. That is, at each update interval of the position of the node, a node changes its direction by $-L$ to $+L$ degrees.

I have selected the realistic model, instead of the Brownian model for example, in order to ensure a degree of correlation between a series of moving iterations.

In the Brownian model, the moving direction of a node at time $(t + 1)$ is generated randomly from a uniform distribution process ranging from 0 to 360 degrees, so it is independent of its previous direction at time t . Thus, in the Brownian model, there may be a possibility where the motion style of a node follows an oscillatory-type trajectory, moving forward and backward, leading to an impractical mobility model.

Miguel Sanchez (misan@disca.upv.es)
Thu, 15 Jul 1999 18:19:03 +0200

Hmmm, Every time we are closer to the idea of modelling mobility in the same way as traffic models. (I've thought before in this way and I think it makes sense). The full picture can be something like viewing non-mobile networks dependent in the traffic dimension and ad-hoc

networks as dependent of the two dimensions of traffic and mobility (at least, this is the way I see it).

In the same way as with traffic models, we can introduce the concept of mobility classes and consider that the full system is built by adding several kind of nodes (likely with different mobility patterns and, depending on the application level also with different traffic patterns).

This is what I understand that Ken is proposing, to build mobility models with a mixture of different patterns for different nodes that, on the other hand and given the example he posted, makes a lot of sense.

Of course I agree that traffic and mobility models are not the only variables in the game.

Chip Elliott (celliott@bbn.com)

Fri, 16 Jul 1999 08:58:24 -0400

>I believe that an ad hoc network is considered to be a *homogeneous* wireless network. That means group of mobile wireless terminals with similar characteristics, regarding technical. Therefore, I believe any simulation scenarios should assume that all terminals present the same behavioral characteristics.

>

I for one do not agree with this. In real ad hoc networks, anything goes. One is allowed to use a variety of radios, power sources, antennas, and comms links. BBN is certainly building ad hoc networks that have very different kinds of nodes and links in them, ranging from fiber optics and gigabit radios with unlimited power, down to handheld devices with omni antennas at the extreme edge of radio range (running at a dozen bits per second). All in the same ad hoc network.

Just to restate earlier thoughts on scenarios, I believe a scenario can be accurately modeled by the following tuple:

- * Pathloss matrix (terrain)
- * Node movements
- * Traffic flows

In short, one models a set of nodes moving across a terrain and exchanging data with each other. If one wants to get fancy, one can model time-varying interference as a combination of time-varying pathloss matrix plus self-interference (which must be measured on a per-transmission or finer basis).

I don't know of any reason why we think a single scenario will accurately capture all the behavior of an ad hoc networking scheme. Therefore it would be best to catalog a standard set of terrains, of movements, and of traffic flows, and of traffic flows, and to run a proposed ad hoc network across this entire range. Sounds tedious but it's really just running a set of perl scripts and collecting the results.

Miguel Sanchez has already made a first catalog of movements. A little more work would give a reasonable set of traffic flows. And then some grinding away with RF propagation tools would give a set of pathloss matrices.

To briefly mount an old soapbox, PLEASE do not do simulations that use free space propagation as the pathloss model. In general these give RIDICULOUSLY MISLEADING results. We'll know that a given protocol suite works ok in the Nevada salt flats, provided there's

no self-interference, but nothing more. Buildings, trees, hills, walls, and elevator shafts do matter!

Justin's Dad (tudball@lis.pitt.edu)

Thu, 15 Jul 1999 12:12:37 -0400 (EDT)

I think that implementing precise individual node mobility models in a simulation is important in order to evaluate and conduct sensitivity analysis of ad-hoc routing performance. However, the question I would like to raise is as follows:

Given that there are a number of good models that describe, at least qualitatively, how we think mobile ad-hoc nodes might move under different scenarios, can we now develop analytical models, which characterize the dynamic behavior of the network topology? Specifically, given two nodes moving according to (for example) the Random Gauss-Markov model proposed by Haas and Liang, and reasonable assumptions governing signal propagation, is there an analytical expression that characterizes the distribution of the time-to-failure for a link between them? Analytical expressions of this type could be very useful---for example in routing, or resource management schemes for ad-hoc networks.

In cell networks the same thing is used to determine system performance measures, e.g. the distribution of the cell residence time is derived and used to evaluate various blocking probabilities, handover times etc.. These are used in system design, also to dynamically allocate channels, and in the future to provide support for probabilistic QoS to mobile users. In fact, many of these models have been based on random/brownian type movement. To solve the same problem in an ad-hoc network introduces a problem---we do not have the benefit of the fixed base station. Therefore, I think the problem is very difficult in general. The problem is even more difficult if you consider the possibility of group movement.

Is there anyone who is attempting to solve this or any similar problems? Does anyone have any insights or feeling relative to value of developing such models? We have developed an analytical model for link availability based on a mobility model that is closely related to the Brownian motion model. Despite the apparent limitations of a brownian type model, it can be an effective model for aggregate node movement in a very large network.

--Bruce McDonald

Ad-Hoc Simulation Model Questions

A. Bruce McDonald (tudball@lis.pitt.edu)

Tue, 9 Nov 1999 13:18:10 -0500 (EST)

I am in the process of designing a model to simulate moderately large (100-1000 nodes) ad-hoc networks with nodes moving at speeds ranging from 10 to 50 kph. My research is focused on network-layer issues (routing), hence, my post to the MANET group! However, as an inevitable consequence of conducting research in wireless networks, it seems that some physical and MAC-layer issues arise...

I am seeking opinions, advice, and references from others on the following modeling issues

- (1) What range of radio transmission ranges (distance) would be realistic and representative of what is available today or might be expected to be available in the near future? Assume a free-space propagation model.
- (2) In the simulation of routing protocols one may conclude that it becomes both computationally excessive, as well as requiring assumptions that may not generalize well to simulate the MAC-layer directly. However, MAC-layer contention is an important component that can affect network-layer performance. Hence, it seems that it should be accounted for in some way that is as general and as computationally efficient as possible.

The question is, does anyone know of any analytical, or empirical models that can be used to estimate the MAC-layer delay that depend on some reasonable set of parameters (eg. node density (mean number of neighbors), mean offered load, etc.) ? For example, Bux derived an analytical expression to model the delay on an Ethernet (802.3) network given a reasonable set of assumptions and parameters.

- (3) What sort of workload models can people suggest? It seems that CBR sources with uniform traffic distributions (random selection of source-destination pairs) is quite common in recent ad-hoc simulations. What other models are being used? Is something like tcplib outdated because it does not reflect current Internet traffic?

Thanks to anyone for any advice on these matters! Please post answers to MANET for everyone's benefit.